Security science and measurement

- Dr. Fred Cohen
  - CEO Fred Cohen & Associates
  - President California Sciences Institute

Outline

- The physics of digital information
- Measurement theory and practice
- Examples of measurements in experiments
- Questions / comments?

Basic concepts and principles

- $C \rightarrow^{m} E$: Cause acts through mechanisms to produce effects
- $t_C < t_E$: Cause precedes effect
- $t_{em} - t_{sm} > 0$: Mechanisms take time to produce effects from causes
- Everything digital has finite granularity (the bit)
  - → Time is a partial ordering
  - → Space is discontinuous, not smooth (& assumptions dangerous)
  - → State space converges with time (while normal space diverges)
  - → The "speed of light" is augmented by computational complexity
  - → Reverse time is ! in possible $C \rightarrow \overline{(E \rightarrow C)}$ (non-unique, large)
- Traces are **not** produced by **transfer**, but by FSM execution
  - We almost never have a complete or equivalent trace
- Consistency and redundancy play heavily in the space
  - Hypothesize, test (for consistency) and refute (if inconsistent)
  - Redundant traces should be consistent!

Outline

- The physics of digital information
- Measurement theory and practice
- Examples of measurements in experiments
- Questions / comments?

Measurement theory

- Metrics options:
  - Ratio metrics (+,-,>, <, 0)
    - Finite granularity → Only available as integers and ratios
    - Very often problematic in the digital space
    - Almost never sensible for security-related measurements
  - Interval metrics (A≤x≤B)
    - Time is essentially always no better than this in digital systems
    - Sequences of bounds may be quite useful
  - Ordinal metrics (partial ordering available)
    - Often available – critical in understanding time and sequences
  - Nominal metrics (make lists, count the lists)
    - Essentially always available
      - How many times did I find “string” in “bigger string”? (once)
    - Often not very meaningful or useful
      - How many viruses were identified last year?
      - How many vulnerabilities were found by the scan?

More on measurements

- Precision - The smallest change in input producing a change in output
- Accuracy - The difference between what is indicated and ground truth
- We often see precision far greater than accuracy
  - 12/17/98 @ 21:22:12.126542 (precise to the nearest microsecond)
  - But it actually happened at or about 2PM Monday (inaccurate)
  - 7 out of 11 (63%) had it (63% more precise than 7 / 11 is accurate)
- Error propagation – how the inaccuracy and imprecision add up
- Assumptions
  - We make lots of them (e.g., it looks like a clock → it is a clock)
  - We need to test assumptions that we make (validate, calibrate)
- Base rates
  - How do we know it's not normal if we don't know normal?
  - We need to measure normal to know what's not normal!

Examples of measurements and calibration

- Measurement: Time it takes to perform an operation
  - Calibration: measure the time through reconstruction
- Measurement: Minimum time granularity (clock resolution)
  - Calibration: measure granularity by trace examination (GCF(Δ))
- Measurement: MAC time sequences vs. claimed actions
  - Calibration: measure MAC time changes by reconstructing acts
- Measurement: Password guessing time from remote locations
  - Calibration: measure password guesses/time from such locations
- The point:
  - We want to measure lots of things
  - But we need to calibrate our equations (and our tools)
  - So we do calibration measurements to identify standards
  - Then we measure against these calibrated standards

Outline

- The physics of digital information
- Measurement theory and practice
- Examples of measurements in experiments
- Questions / comments?

Example: Detecting insiders breaking rules

- Idea: Insiders turning break rules of certain types
- CERT reports for several years on insiders tell us things like:
  - X% of “bad” insiders who were caught deleted files
  - Y% of “bad” insiders who were caught used another user's UID
  - Z% of “bad” insiders who were caught were male
  - Etc.
- However, **no base rate data** was apparently collected or analyzed
  - What % of ALL insiders deleted files?
  - What % of ALL insiders used another users UID?
  - What % of ALL insiders were male?
- Without the base rates, we cannot differentiate “indicators” from “random” or assess the utility of the measurement
  - Why do we get so many false positives in IADRS? No base rates?
  - Why is it so easy to avoid detection? Too little time to investigate because of the lack of base rates?

Approach: Look for inconsistencies in traces

Example: Check CAC / badge / computer timestamps for consistency

Assumptions for timestamp consistency analysis (ongoing / expanding)

- entry/exit of areas is always recorded
- you can't swipe a card at the entry/exit and then not enter/exit
- you can't enter/exit without swiping
- entry and exit use the same clock
- we know when one area is inside another
- we have complete knowledge of person/card/... identities
- same-person, same-card
- one person per card
- recorded commands require the presence of a person at a terminal
- terminals and areas do not move
- minimum travel times do not change
- first entry must precede first exit
- person who never entered is outside

Testing those hypotheses by measurement

- entry/exit of areas is always recorded (red teaming / log examination)
- you can't swipe a card at the entry/exit and then not enter/exit (try it)
- you can't enter/exit without swiping (red teaming / log examination)
- entry and exit use the same clock (log examination / try it)
- we know when one area is inside another (physical examination)
- we have complete knowledge of person/card/... identities
- same-person, same-card (physical examination)
- one person per card (physical examination)
- recorded commands require the presence of a person at a terminal
- terminals and areas do not move (we know it isn't so because of ships)
- minimum travel times do not change (red teaming / log examination)
- first entry must precede first exit (log examination)
- person who never entered is outside (red teaming / log examination)

There are many more hypotheses

- Measurement must be applied to each based on the needs of the use
  - The measurement (experimental) process must be done properly
  - The things measured must reflect the phenomena of interest
  - The precision and accuracy of measurement must reflect the need
- Example measurement – travel time (physical space)
  - Measure travel time from location l_1 to location l_2
  - Repeated experiments looking for minimum times
  - Augment with theoretical analysis (min of each link in the graph)
  - Augment with margin of error to desired likelihood
  - Compare to recorded sequences of timestamps in records
  - Investigate any discrepancies till resolved

Example measurement – people in places

- People who appear inside without entering
  - Hypothetically, “secure areas” have “controlled” entry
  - Hypothetically, to enter you must “badge in”
  - Realistically, we have:
    - Vouching
    - Tailgating
    - Jumping the fence – likely highly discouraged
    - Alternative entry modes (fire, ambulance, guard checks, etc.)
- Question: Can we use presence inconsistencies? What kinds?
  - Measure presence inconsistencies by trace analysis
  - Check out each inconsistency for true positives
  - Toss out true positives and find root cause for false positives
  - Change the rules of the game
    - No vouching, technical tailgate controls, enter exceptions for emergency modes, etc.
  - Select for low base rate phenomena

Example measurement – MAC times

- MAC:= Modify / Access / Create – timestamps in files/directories
- Assumption: Some are invalid sequences (e.g. C > A, C > M)
- These assumptions may be wrong
  - C is not necessarily create – it is directory change time on Unix
  - Timestamps may have different resolutions
  - Different commands may have different effects (mv, cp, tar, etc.)
  - System calls may alter one and not the other (settime)
  - Physical alteration of media may effect times
  - Different device drivers / file systems may produce different times
- To find out we have to test different mechanisms in different situations
  - A generic test won't necessarily be right – nor will assumptions
  - Measure by testing in situ – with actual commands from system
  - Self-calibrate tools by testing each time
  - Leads to situation-specific $C \rightarrow^{m} E$
  - Analysis is then based on situation specifics and not generics

MAC time self-calibration forensic tool

- Tool does inconsistency analysis between hypotheticals and traces
  - Look at traces to identify possible causes of effects
    - e.g., look at shell logs for commands that could have copied a file to a remote server
  - For each candidate cause, test in situ – e.g.,
    - Boot a forensically sound image of the machine and test each command in a simulated external environment
    - For each command from the shell logs, examine the results of running that command and examine the resulting metadata
  - ∀ inconsistent trace results, consider it refutation
  - ∀ consistent / indeterminate results, recurse back the causal chain
- We then have candidates for causes of the effects, but only candidates
  - Candidates are consistent with the traces BUT $\overline{(E \rightarrow C)}$

Court case: US v. Bayley, et. al.

- Defendant Fuhs accused of conspiracy to commit fraud (along with the other Enron defendants) and lying to investigators
- Lying to investigators was the denial that he participated in the fraud
- The case for fraud was based on traces of a file received in email
  - Claim: Fuhs received the file, added a key phrase, and sent it back
  - Key point: If he did, then he was a knowing participant in the fraud
- The evidence was in the form of a single file found on a file server
  - All the other evidence was stored in the WTC basement
  - The time frames were critical (w/in an hour several years earlier)
  - The file was a Microsoft Word document
    - Which (was) an Object Linking and Embedding (OLE) file
- OLE files contain timestamps for different “objects” they contain
  - 2 creation timestamps each
  - These timestamps are undocumented at to how they came to be
  - Most tools ignore the 2nd one, which is usually identical to the 1st
  - But not in this case

Court case: US v. Bayley, et. al. (cont)

- The timestamps were different in this file
  - The 2nd one was offset by 5 seconds from the 1st one
  - But what does this mean?
- Hypothesis: One is creation, the other modification
  - If so, Fuhs had only 20 seconds of editing and could not have done what was claimed he did
- Hypothesis: They should never differ
  - If so, the file is a forgery, and someone forgot to fix the 2nd date
- Hypothesis: We can speculate about lots of other hypotheses
- Some other issues:
  - The file was saved on a file server in Houston in a Fuhs directory
  - It was the only copy of the file at issues found
  - Other earlier generations were found elsewhere, but the record was incomplete
- We decided to try reconstruction to try to determine what this and other metadata in the OLD file meant in terms of the case at hand

Court case: US v. Bayley, et. al. (cont)

- The reconstruction background
  - The file was apparently created from an email sent to Fuhs
    - Records were unrevealing as to which email
  - The company used Microsoft / Exchange server / Mail client
  - The file was retrieved from a server where it was saved apparently by Fuhs upon or after receipt in Texas (Fuhs was in New York)
  - These leave different timestamps in the file base don how things are done and the different patch versions in place at the time
  - No records of the patch versions in place were available
- The reconstruction approach
  - Create VMs to model the exchange server, network, etc.
  - Create a Windows version based on the metadata from files
  - Use Samba to emulate Widows file shares at different locations
  - Reboot, do email exchanges, save the file in different ways
  - Stop system, examine metadata, rule out patch level / or not
  - Reboot system, load the next patch in the series, redo it all
  - Loop till last patch available before operative date

Court case: US v. Bayley, et. al. (cont)

- Results of the reconstructions (experiments)
  - One and only one patch level produced the right metadata
  - Different ways of saving the file produced different timestamp data
    - The offset dates are different from different methods
      - Offset from Jan 1, 1400, Offset from Feb 1, 1962, etc.
    - The differential between the 1st and 2nd timestamps was only found in one class of file save methods
- Between the various combinations of results, we found:
  - At the particular patch level
  - With the particular "Save-As" method
    - Keyboard shortcuts are different from menu selections
  - In the particular location saved (network is different than local)
  - We reproduced the time differential between the timstamps
- The 1st is from the computer, the 2nd from the filesystem
  - Hypotheses refuted – result indeterminate in terms of the case
  - This cannot be the basis for claims of time spent editing

Court case: US v. Bayley, et. al. (cont)

- But there's more...
  - The file had “last 10” data – So what is “Last 10” data?
  - Many claim it is a record of the last 10 users who edited a file
  - Fuhs was indicated as 8 of the Last 10 data entries (I think)
  - Prosecution expert claims that this shows Fuhs edited the document over a long time frame
  - But there is also a record of edit time – and it was 0!
  - But edit time is set to 0 when a “Save-As” is done – which my reconstruction showed was done
- So Fuhs must have edited the file and done a Save-As – right?
  - Wrong!
  - Last 10 was not documented as to actual function
  - The commercial software claiming to retrieve it disclaims reliability and will not answer questions about what it does or how it works
  - In a reconstruction we found that ALL unused Last 10 slots were replaced by the current UID the 1st time a file was received and a if a “Save-As” was done immediately

The case?

- I testified as the last witness – surrebuttal
  - There were 7 or 8 defendants in this particular case
  - All but 1 were convicted at trial
- Fuhs was convicted on both counts - GUILTY
  - Fuhs started his long jail term
  - But on appeal the case was reversed with a directed verdict
    - NOT GUILTY
- Fuhs was released after serving a few months in jail
- BECAUSE the digital evidence was **not determinative**

- And the science of digital forensics continued to move forward...
  - In your dreams...

Outline

- The physics of digital information
- Measurement theory and practice
- Examples of measurements in experiments
- Questions / comments?

The truth of information security science

- It is not advancing very rapidly – but science rarely does
  - No identifiable funding for basic science
    - Lots of things called science
    - Rarely any real science in them
      - No underlying notions like:
        - $C \rightarrow^{m} E$: Cause via mechanisms produce effects
        - $t_C < t_E$: Cause before effect, $t_{em} - t_{sm} > 0$: and takes time
      - No requirement to use existing theory as a foundation
  - Widespread lack of consensus in the "scientific" community
  - No common language (although some progress has been made)
  - No repetition in experiments
  - Lots of human experimentation WITHOUT proper IRB approval
- Security science is hard, expensive, slow, complex, poorly supported
  - EXCEPT at DoE (which has done good research for a long time)
- Why should information security science be any different?

Thank You

http://calsci.org/ - calsci at calsci.org

http://all.net/ - fc at all.net